

Structural parameterizations for boxicity

Henning Bruhn, Morgan Chopin, Felix Joos and Oliver Schaudt

Abstract

The boxicity of a graph G is the least integer d such that G has an intersection model of axis-aligned d -dimensional boxes. BOXICITY, the problem of deciding whether a given graph G has boxicity at most d , is NP-complete for every fixed $d \geq 2$. We show that BOXICITY is fixed-parameter tractable when parameterized by the cluster vertex deletion number of the input graph. This generalizes the result of Adiga et al., that BOXICITY is fixed-parameter tractable in the vertex cover number.

Moreover, we show that BOXICITY admits an additive 1-approximation when parameterized by the pathwidth of the input graph.

Finally, we provide evidence in favor of a conjecture of Adiga et al. that BOXICITY remains NP-complete when parameterized by the treewidth.

1 Introduction

Every graph G can be represented as an intersection graph of axis-aligned boxes in \mathbb{R}^d , provided d is large enough. The *boxicity* of G , introduced by Roberts [18], is the smallest dimension d for which this is possible. We denote the corresponding decision problem by BOXICITY: given G and $d \in \mathbb{N}$, determine whether G has boxicity at most d .

Boxicity has received a fair amount of attention. This is partially due to the wider context of graph representations, but also because graphs of low boxicity are interesting from an algorithmic point of view. While many hard problems remain so for graphs of bounded boxicity, some become solvable in polynomial time, notably max-weighted clique (as observed by Spinrad [20, p. 36]).

Cozzens [9] showed that BOXICITY is NP-complete. To cope with this hardness result, several authors [1, 3, 14] studied the parameterized complexity of BOXICITY. Since the problem remains NP-complete for constant $d \geq 2$ (Yannakakis [22] and Kratochvíl [16]), boxicity itself is ruled out as parameter. Instead more structural parameters have been considered. Our work follows this line. We prove:

Theorem 1. *BOXICITY is fixed-parameter tractable when parameterized by cluster vertex deletion number.*

The *cluster vertex deletion* number is the minimum number of vertices that have to be deleted to get a disjoint union of complete graphs or *cluster graph*. As discussed by Doucha and Kratochvíl [11] cluster vertex deletion is an intermediate parameterization between vertex cover and cliquewidth. A *d -box representation* of a graph G is a representation of G as intersection graph of axis-aligned boxes in \mathbb{R}^d .

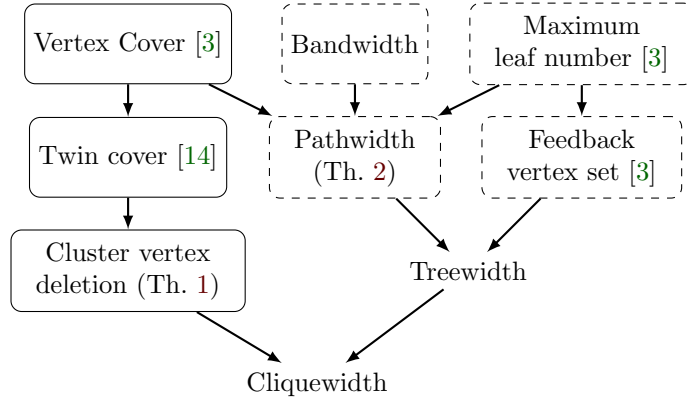


Figure 1: Navigation map through our parameterized complexity results for BOXICITY. An arc from a parameter k_2 to a parameter k_1 means that there exists some function h such that $k_1 \leq h(k_2)$. A rectangle means fixed-parameter tractability for this parameter and a dashed rectangle means an approximation algorithm with running time $f(k) \cdot n^{O(1)}$ is known.

Theorem 2. *Finding a d -box representation of G such that $d \leq \text{box}(G) + 1$ can be done in $f(\text{pw}(G)) \cdot |V(G)|$ time where $\text{pw}(G)$ is the pathwidth of G .*

A natural parameter for BOXICITY is the treewidth $\text{tw}(G)$ of a graph G , in particular as Chandran and Sivadasan [7] proved that $\text{box}(G) \leq \text{tw}(G) + 2$. However, Adiga, Chitnis and Saurabh [3] conjecture that BOXICITY is NP-complete on graphs of bounded treewidth. Our last result provides evidence in favor of this conjecture. For this, we mention the observation of Roberts [18] that a graph G has boxicity d if and only if G can be expressed as the intersection of d interval graphs.

Theorem 3. *There is an infinite family of graphs G of boxicity 2 and bandwidth $\mathcal{O}(1)$ such that, among any pair of interval graphs whose intersection is G , at least one has treewidth $\Omega(|V(G)|)$.*

Why do we see the result as evidence? An algorithm solving BOXICITY on graphs of bounded treewidth (or even stronger, of bounded bandwidth) is likely to exploit the local structure of the graph in order to make dynamic programming work. Yet, Theorem 3 implies that this locality may be lost in some dimensions, which constitutes a serious obstacle for any dynamic programming based approach. We discuss this in more detail in Section 5.

Figure 1 summarizes previously known parameterized complexity results on boxicity along with those obtained in this article. Adiga et al. [3] initiated this line of research when they parameterized BOXICITY by the minimal size k of a *vertex cover* in order to give an $2^{O(2^k k^2)} \cdot n$ -time algorithm, where n denotes the number of vertices of the input graph, as usual. Adiga et al. also described an approximation algorithm that, in time $2^{O(k^2 \log k)} \cdot n$, returns a box representation of at most $\text{box}(G) + 1$ dimensions. Both results were extended by Ganian [14] to the less restrictive parameter *twin cover*. Our Theorem 1 includes Ganian's.

Other structural parameters that were considered by Adiga et al. for parameterized approximation algorithms are the size of a *feedback vertex set* – the minimum number of vertices that need to be deleted to obtain a forest – and *maximum leaf number* – the maximum number of leaves in a spanning tree of the graph. They proved that finding a d -box representation of a graph G such that $d \leq 2 \text{box}(G) + 2$ (resp. $d \leq \text{box}(G) + 2$) can be done in $f(k) \cdot |V(G)|^{O(1)}$ time (resp. $2^{O(k^3 \log k)} \cdot |V(G)|$ time) where k is the size of a feedback vertex set (resp. maximum leaf number). In [1], Adiga, Babu, and Chandran generalized these approximation algorithms to parameters of the type “distance to \mathcal{C} ”, where \mathcal{C} is any graph class of bounded boxicity. More precisely, the parameter measures the minimum number of vertices whose deletion results in a graph that belongs to \mathcal{C} .

The algorithm of Theorem 2 generalizes the approximation algorithm for the parameter vertex cover number, and improves the guarantee bound of the approximation algorithm for the parameter maximum leaf number.

There is merit in studying approximation algorithms from a parameterized perspective: not only is BOXICITY NP-complete, but the associated minimization problem cannot be approximated in polynomial time within a factor of $n^{\frac{1}{2}-\varepsilon}$ for any $\varepsilon > 0$ even when the input is restricted to bipartite, co-bipartite or split graphs (provided $\text{NP} \neq \text{ZPP}$). This is a result due to Adiga, Bhowmick and Chandran [2]. There is, however, an approximation algorithm with factor $o(n)$ for general graphs; see Adiga et al. [1].

While Roberts [18] was the first to study the boxicity parameter, he was hardly the first to consider box representations of graphs. Already in 1948 Bielecki [5] asked, here phrased in modern terminology, whether triangle-free graphs of boxicity ≤ 2 had bounded chromatic number. This was answered affirmatively by Asplund and Grünbaum in [4]. Kostochka [15] treats this question in a much more general setting.

Following Roberts who proved that $\text{box}(G) \leq \frac{n}{2}$, other authors obtained bounds for boxicity. Esperet [13], for instance, showed that $\text{box}(G) \leq \Delta(G)^2 + 2$, while Scheinerman [19] established that every outerplanar graph has boxicity at most two. This, in turn, was extended by Thomassen [21], who showed that planar graphs have boxicity at most three.

In the next section, we will give formal definitions of the necessary concepts for this article. We prove our main results in Sections 3–5. Finally, we discuss the impact and limitations of our results in Section 6, where we also outline some future directions for research.

2 Preliminaries

Graph terminology. We follow the notation of Diestel [10], where also all basic definitions concerning graphs may be found.

Let X be some finite set. With a slight abuse of notation, we consider a collection $I = ([\ell_v, r_v])_{v \in X}$ of closed intervals in the real line to be an *interval graph*: I has vertex set X , and two of its vertices u and v are adjacent if and only if the corresponding intervals $[\ell_u, r_u]$ and $[\ell_v, r_v]$ intersect. By perturbing the endpoints of the intervals we can ensure that no two intervals have a common endpoint, and that for every interval the left endpoint is distinct from the right

endpoint. We always tacitly assume the intervals to be of that form. Fig. 2 shows the family of forbidden subgraphs for the class of interval graphs.

The *bandwidth* of a graph G , say with vertex set $V(G) = \{v_1, v_2, \dots, v_n\}$, is the least number k for which the vertices of G can be labeled with distinct integers $\ell(v_i)$ such that $k = \max\{|\ell(v_i) - \ell(v_j)| : v_i v_j \in E\}$. Equivalently, it is the least integer k for which the vertices of G can be placed at distinct integer points on the real line such that the length of the longest edge is at most k . We denote the bandwidth of a graph G by $\text{bw}(G)$.

We define a *path decomposition* of a graph G as a set $\mathcal{W} = \{W_1, \dots, W_t\}$ of subsets of $V(G)$ called *bags* such that the following conditions are met.

1. $\bigcup_{i=1}^t W_i = V(G)$.
2. For each $uv \in E(G)$, there is an $i \in \{1, \dots, t\}$ such that $u, v \in W_i$.
3. For each $v \in V(G)$, if $v \in W_i \cap W_j$ for some $i, j \in \{1, \dots, t\}$, then $v \in W_k$ with $i \leq k \leq j$.

The *width* of a path decomposition is $\max_i |W_i| - 1$. The *pathwidth* $\text{pw}(G)$ of a graph G is the minimum width over all possible path decompositions of G . Equivalently, $\text{pw}(G)$ is the minimum size of the largest clique of any interval supergraph of G , minus 1.




Figure 2: Forbidden induced subgraphs for interval graphs; the dashed paths may have any length ≥ 1

The *treewidth* of a graph G , denoted $\text{tw}(G)$, is the minimum size of the largest clique of any chordal supergraph of G , minus 1. For the purpose of our paper it is important to remark that for every graph G we have $\text{tw}(G) \leq \text{pw}(G) \leq \text{bw}(G)$.

Parameterized complexity. A decision problem parameterized by a problem-specific parameter k is called *fixed-parameter tractable* if there exists an algorithm that solves it in time $f(k) \cdot n^{O(1)}$, where n is the instance size. The function f is typically super-polynomial and only depends on k . One of the main tools to design such algorithms is the *kernelization* technique. A kernelization algorithm transforms in polynomial time an instance I of a given problem parameterized by k into an equivalent instance I' of the same problem parameterized by $k' \leq k$ such that the size of I' is bounded by $g(k)$ for some computable function g . The instance I' is called a *kernel* of size $g(k)$. The following folklore result is well known.

Theorem 4. *A parameterized problem P is fixed-parameter tractable if and only if P has a kernel.*

In the remainder of this paper, the kernel size is expressed in terms of the number of vertices.

For more background on parameterized complexity the reader is referred to Downey and Fellows [12].

Problem definition. We call an *axis-aligned d -dimensional box* (or *d -box*) a cartesian product of d closed real intervals. A *d -box representation* of a graph G is a mapping that maps every vertex $v \in V(G)$ to a d -box B_v such that two vertices $u, v \in V(G)$ are adjacent if and only if their associated boxes have a non-empty intersection. The *boxicity* of G , denoted by $\text{box}(G)$, is the minimum integer d such that G admits a d -box representation. We consider the following problem.

BOXICITY

Input: A graph G and an integer d .

Question: Is $\text{box}(G) \leq d$?

Given a d -box representation of G , we denote by $[\ell_i(v), r_i(v)]$ the interval representing v in the i -th dimension.

Throughout the article, we make frequent use of the reformulation of boxicity in terms of interval graphs:

Theorem 5 (Roberts [18]). *The boxicity of a graph G is equal to the smallest integer d so that G can be expressed as the intersection of d interval graphs.*

3 Proof of Theorem 1

Theorem 1 follows immediately from the following lemma:

Lemma 6. BOXICITY admits a kernel of at most $k^{2^{O(k)}}$ vertices, where k is the cluster vertex deletion number of the input graph.

In the course of this section, we present a sequence of lemmas in order to prepare the proof of our main lemma above.

Two adjacent vertices u, v in a graph G are *true twins* if u and v have the same neighbourhoods in $G - \{u, v\}$. As observed by Ganian [14], deleting one of two true twins does not change the boxicity.

Lemma 7. *Let u, v be true twins of a graph G . Then $\text{box}(G) = \text{box}(G - u)$.*

We remark, without proof, that there is also a reduction for false twins (those that are non-adjacent): if there are at least three of them, then one may be deleted without changing the boxicity. We will not, however, make use of this observation.

Recall that a *cluster graph* is the disjoint union of complete graphs, called *clusters*. In what follows, we implicitly identify a cluster with its vertex set.

Let $G - X$ be a cluster graph for some $X \subseteq V(G)$. We call two clusters C, C' of $G - X$ *equivalent* if there is a bijection $C \rightarrow C'$, $v \mapsto v'$, such that $N_G(v) \cap X = N_G(v') \cap X$. Observe that, if $G - X$ has no true twins, then

two clusters C and C' are equivalent if and only if $\{N_G(u) \cap X : u \in C\} = \{N_G(v) \cap X : v \in C'\}$.

Lemma 8. *Let G be a graph without true twins, and let X be a set of k vertices so that $G - X$ is a cluster graph. Then every cluster in $G - X$ contains at most 2^k vertices.*

Proof. Consider a cluster C of G . Then the number of sets $N_G(v) \cap X$, $v \in C$, is bounded by 2^k . As G has no true twins, no two vertices in C may have the same neighbourhood in X , which implies that $|C| \leq 2^k$. \square

We also need the following result.

Theorem 9 (Chandra and Sivadasan [7]). *It holds that $\text{box}(G) \leq \text{tw}(G) + 2$ for any graph G .*

In particular, $\text{box}(G) \leq \text{pw}(G) + 2$ for any graph G .

Lemma 10. *Let G be a graph without true twins, and let X be a set of k vertices so that $G - X$ is a cluster graph. Moreover, let \mathcal{D} be an equivalence class of clusters with $|\mathcal{D}| \geq 2(2k+2)^{2^{k+1}(2^k+k+1)}$. For every $C^* \in \mathcal{D}$, $\text{box}(G) = \text{box}(G - C^*)$.*

Proof. As deleting vertices may only decrease the boxicity, it suffices to prove that $\text{box}(G) \leq \text{box}(G - C^*)$.

Set $H = G - C^*$, $d = \text{box}(H)$, $k = |X|$ and $\mathcal{C} = \mathcal{D} \setminus \{C^*\}$. We claim that

$$d = \text{box}(H) \leq 2^k + k + 1. \quad (1)$$

Indeed, define a path decomposition with a bag W_C for every cluster C of $H - X$ such that $W_C = X \cup C$. This gives a path decomposition of H with width at most $k + 2^k - 1$, by Lemma 8. Theorem 9 now implies (1).

For the sake of simplicity, let us introduce the following notions. Fix a d -box representation of H . The set of *corners* of a box of a vertex is the cartesian product $\times_{i=1}^d \{\ell_i(v), r_i(v)\}$. By rescaling every dimension (compare Lemma 12), we can ensure that every endpoint of an interval of a vertex in X lies in $\{1, 2, \dots, 2k\}$. Thus every corner of a box of X lies in the grid $\{1, 2, \dots, 2k\}^d$. We may moreover assume that every other box of H is contained in $[0, 2k+1]^d$. Points of $\{0, 1, \dots, 2k+1\}^d$ we call *grid points*, and any set $[z_1, z_1+1] \times \dots \times [z_d, z_d+1]$, where $z_i \in \{0, \dots, 2k\}$, is a *grid cell*. In each dimension i we say that the grid induces the *grid intervals* $[0, 1], [1, 2], \dots, [2k, 2k+1]$. A box of a vertex in $H - X$ is a *cluster box*.

By perturbing the boxes slightly we may always assume that

$$\begin{aligned} &\text{if } s \text{ is a corner of a cluster box of a cluster } C \text{ of } H - X, \text{ and if} \\ &t \text{ is a corner of the box of any vertex } z \in V(H - C) \text{ then } s_i \neq t_i \\ &\text{for all dimensions } i = 1, \dots, d. \end{aligned} \quad (2)$$

Moreover, we may assume that any corner of a cluster box lies in the interior of a grid cell. A cluster box that does not contain any grid point we call a *thin box*.

We concentrate on *thin* clusters, that is, clusters that consist of thin boxes only. We claim that

$$\text{at least } (2k+2)^{2^{k+1}(2^k+k+1)} \text{ clusters in } \mathcal{C} \text{ are thin.} \quad (3)$$




Figure 3: Boxes A, B are in the same position, as are C and D ; F is not thin.

To prove this claim, observe that no grid point lies in a cluster box of two different clusters as then two vertices in distinct clusters would be adjacent. Thus, there is at most one cluster per grid point so that one of its cluster boxes contains the grid point. As, by (1), there are $(2k+2)^d \leq (2k+2)^{2^k+2k+1}$ grid points, it follows that \mathcal{C} has at least $|\mathcal{C}| - (2k+2)^{2^k+2k+1} \geq (2k+2)^{2^{k+1}(2^k+k+1)}$ thin clusters.

We say that two cluster boxes B and B' are in the *same position* if every grid cell containing a corner of B also contains a corner of B' and vice versa (see Figure 3). Note that if two vertices $v, v' \in V(H) - X$ have boxes in the same position then $N_H(v) \cap X = N_H(v') \cap X$. (Here we use the fact that cluster boxes have their corner strictly in the interior of grid cells.)

For every cluster $C \in \mathcal{C}$ we fix a point $p(C)$ that lies in every cluster box of C : such a point exists by the Helly property for boxes in \mathbb{R}^d . We claim that, using this Helly point, we can modify our box representation of H so that

$$\begin{aligned} &\text{for all thin clusters } C \in \mathcal{C} \text{ and for each dimension } i \in \{1, \dots, d\} \\ &\text{holds the following: if } p(C) \text{ and a corner } t \text{ of a box of } C \text{ lie in} \\ &\text{the same grid interval in dimension } i, \text{ that is, if there is a } j \text{ so} \\ &\text{that } p_i(C), t_i \in [j, j+1], \text{ then } t_i = p_i(C). \end{aligned} \quad (4)$$

To achieve (4), we proceed as follows. Let v be a vertex of any thin cluster $C \in \mathcal{C}$. Consider a dimension i where $\ell_i(v)$ or $r_i(v)$ lie in the same grid interval as $p_i(C)$. Note that $\ell_i(v) \leq p_i(C) \leq r_i(v)$. In dimension i , we shrink the box of v in the following way: if $\ell_i(v)$ lies in the same grid interval as $p_i(C)$, we replace $\ell_i(v)$ by $p_i(C)$. Similarly, if $r_i(v)$ lies in the same grid interval as $p_i(C)$, we replace $r_i(v)$ by $p_i(C)$. This procedure is illustrated in Fig. 4.

Since by shrinking a box we may only lose edges of the corresponding graph, it suffices to show that every edge is still present. Since the new box of v still contains $p(C)$, the vertex v is still adjacent to every other vertex in C . As we change the box of v only within a grid interval, the old and the new box of v are in the same position. Thus, we do not lose any edge from v to X . Performing this transformation iteratively for every box of C in every dimension, and for every thin cluster $C \in \mathcal{C}$, we obtain a box representation of H satisfying (4).




Figure 4: Shrinking the boxes

Next, we claim that

there is a pair of distinct thin clusters $C, C' \in \mathcal{C}$ such that for every $v \in C$ and $v' \in C'$ with $N_H(v) \cap X = N_H(v') \cap X$, the boxes of v and v' are in the same position. (5)

Note that, as C and C' are equivalent, there is indeed a bijection between the vertices of C and C' that maps a vertex v to $v' \in C'$ with $N_H(v) \cap X = N_H(v') \cap X$.

Observe that for the endpoints $\ell_i(v), r_i(v)$ of the interval representing a vertex $v \in V(H)$ in the i -th dimension, there are at most $(2k+1)^2$ many choices to select the grid intervals they lie in. Thus, any set of thin boxes, pairwise not in the same position, has size at most $(2k+1)^{2d}$. Because G is devoid of true twins, no cluster has two vertices whose boxes are in the same position.

Recall that every cluster has at most 2^k vertices. Thus, among any choice of more than $(2k+1)^{2d \cdot 2^k}$ thin clusters there are two thin clusters satisfying (5). As $(2k+1)^{2d \cdot 2^k} \leq (2k+1)^{2(2^k+k+1) \cdot 2^k}$, by (1), and since \mathcal{C} contains at least $(2k+2)^{2^{k+1}(2^k+k+1)}$ thin clusters, by (3), the claim follows.

Consider clusters C, C' as in (5). We now embed the deleted cluster C^* in the box representation of $H = G - C^*$. For this, choose $\epsilon > 0$ small enough so that

for all $v \in C$ and $w \in V(H - C)$ and all dimensions i it holds that $|s_i - t_i| > \epsilon$, when s is a corner of the box of v and t is a corner of the box of w . (6)

(If such an ϵ does not exist, we may again perturb the box representation slightly so as to guarantee (2) while keeping (4).)

Define $q \in \mathbb{R}^d$ by setting

$$q_i = \begin{cases} 1 & \text{if } p_i(C) < p_i(C') \\ -1 & \text{if } p_i(C) > p_i(C') \\ 0 & \text{if } p_i(C) = p_i(C'). \end{cases}$$

Let $v \mapsto v^*$ be the bijection between C and C^* with $N_G(v) \cap X = N_G(v^*) \cap X$. We define a box for every $v^* \in C^*$ by taking a copy of the box of v and shifting its coordinates by the vector $\epsilon \cdot q$, that is, for every dimension i we set

$$\ell_i(v^*) = \ell_i(v) + \epsilon q_i \text{ and } r_i(v^*) = r_i(v) + \epsilon q_i$$

Note that, by choice of ϵ , the box of v^* and the box of v are in the same position.

Let \tilde{G} be the graph defined by this new box representation. We claim that $\tilde{G} = G$, which then finishes the proof of the lemma.

To prove this, we first note that we only added edges between vertices in C^* and H , while all other adjacencies remain unchanged. Next, as $p(C) + \epsilon q$ is a point that lies in every box of C^* , it follows that $\tilde{G}[C^*]$ is a complete graph. Moreover, by choice of ϵ , we have

$$N_{\tilde{G}}(v^*) \setminus (C \cup C^*) = N_G(v) \setminus (C \cup C^*)$$

for any $v \in C$. In particular, $N_{\tilde{G}}(v^*) \cap C' = \emptyset$. It remains to show that also $N_{\tilde{G}}(v^*) \cap C = \emptyset$.

For this, let $w^* \in C^*$ and $v \in C$ be arbitrary, where we allow that $v = w$. Let us show that the boxes of v and w^* do not intersect.

Since v and w' are nonadjacent in H , there is a dimension i such that either $r_i(v) < \ell_i(w')$ or $r_i(w') < \ell_i(v)$. By symmetry, we may assume $r_i(v) < \ell_i(w')$. Let I be the grid interval such that $r_i(v) \in I$. If $\ell_i(w') \notin I$, then $r_i(v) < \ell_i(w^*)$, since by our construction $\ell_i(w^*)$ is in the same grid interval as $\ell_i(w')$. This means that the boxes of v and w^* do not intersect. Thus, we may assume that $\ell_i(w^*) \in I$. As v and w are in the same cluster and thus adjacent, it follows that $\ell_i(w) \leq r_i(v)$, which implies that $p_i(C) \in [\ell_i(w), r_i(v)] \subseteq I$. Now, (4) implies that $r_i(v) = p_i(C) = \ell_i(w)$.

Since $p_i(C) = r_i(v) < \ell_i(w')$, it follows that $p_i(C) < p_i(C')$. Thus, $r_i(v) = \ell_i(w) < \ell_i(w) + \epsilon = \ell_i(w^*)$. Consequently, the boxes of v and w^* do not intersect. This completes the proof. \square

We can now prove the main lemma of this section.

Proof of Lemma 6. Let $I = (G = (V, E), b)$ be an instance of BOXICITY with cluster vertex deletion number k . We first compute a set X of size $|X| \leq 3k$, so that $G - X$ is a cluster graph. To this end we use the fact that a graph is a cluster graph if and only if it does not contain a P_3 as an induced subgraph. Start with $X = \emptyset$. If G contains three vertices v_1, v_2, v_3 that induce a P_3 in G then add these three vertices to X . Reiterate the process on $G \setminus \{v_1, v_2, v_3\}$ until no more induced P_3 is found. Clearly any optimal solution needs to delete at least one vertex in an induced P_3 . Thus $|X| \leq 3k$.

Next, we iteratively remove one twin of any pair of true twins from G until the graph becomes free of true twins. By Lemma 7, this does not change the boxicity of G . In the next step, we divide the clusters of $G - X$ into their equivalence classes which can be done in polynomial time. Then we delete clusters from every equivalence class until each equivalence class has at most $2(2k+2)^{2^{k+1}(2^k+k+1)}$ members. Since every cluster has size at most $2^{|X|} \leq 2^{3k}$, by Lemma 8, the resulting graph H has size at most $k^{2^{O(k)}}$. Moreover, Lemma 10 shows that $\text{box}(H) = \text{box}(G)$. This completes the proof. \square

4 Proof of Theorem 2

Bounded pathwidth suggest a dynamic programming approach, and this is precisely what we do. There is a hitch, though. The standard approach would be

to solve the BOXICITY problem on one bag after another of the path decomposition, so that the local solutions can be combined to a global one. BOXICITY, however, does not permit this: as we are constructing the box representation of the graph, we may have to completely rearrange the previous boxes to add a new one.

Thus, the key issue is to force the problem to become “localized”. To this end, we introduce a special interval graph I^* that reflects the path structure of the graph: two vertices are adjacent if and only if they appear in the same bag of the path decomposition. Doing so, we can safely compute local box representations of the subgraphs induced by the bags without paying attention to how these representations overlap. Indeed, the interval graph I^* gets rid of any unwanted adjacency.

Having sketched the idea, we now give the formal description of the algorithm. We say that two interval graphs $I = ([\ell_v, r_v])_{v \in V}$ and $I' = ([\ell'_v, r'_v])_{v' \in V'}$ are *consistent* if the order of the interval endpoints of the common vertices is the same, that is, if for all $u, v \in V \cap V'$

$$\begin{aligned} \ell_u \leq \ell_v &\Leftrightarrow \ell'_u \leq \ell'_v; & r_u \leq r_v &\Leftrightarrow r'_u \leq r'_v; \\ \ell_u \leq r_v &\Leftrightarrow \ell'_u \leq r'_v; & \text{and } r_u \leq \ell_v &\Leftrightarrow r'_u \leq \ell'_v. \end{aligned}$$

In particular, if $V' \subseteq V$ then I' is an induced subgraph of I . When we consider tuples (I_1, \dots, I_d) and (I'_1, \dots, I'_d) of interval graphs we say that they are *consistent* if I_i is consistent with I'_i for $i = 1, \dots, d$.

Lemma 11. *If $I = ([\ell_v, r_v])_{v \in V}$ and $I' = ([\ell'_v, r'_v])_{v' \in V'}$ are two consistent interval graphs, then there is an interval graph J on $V \cup V'$ that is consistent with both I and I' . In particular, any edge in $E(J) \setminus (E(I) \cup E(I'))$ has one endvertex in $V \setminus V'$ and the other in $V' \setminus V$.*

Proof. Pick $v^* \in V' \setminus V$ and apply induction to the pair of consistent interval graphs I and $I' - v^*$ in order to obtain an interval graph J' on $(V \cup V') \setminus \{v^*\}$ that is consistent with I and with I' . Let $J' = ([\tilde{\ell}_v, \tilde{r}_v])_{v \in (V \cup V') \setminus \{v^*\}}$. Among the interval endpoints of $I' - v^*$, that is in $\bigcup_{v' \in V' \setminus \{v^*\}} \{\ell'_{v'}, r'_{v'}\}$, we pick the two consecutive endpoints p, q for which $p \leq \ell'_{v^*} \leq q$. Notice that the endpoint p (resp. q) might not exist; in such case we simply set $p = \ell'_{v^*}$ (resp. $q = \ell'_{v^*}$). Let p' (resp. q') be the corresponding endpoint of p (resp. q) in J' . Then we put $\tilde{\ell}_{v^*} = \frac{1}{2}(p' + q')$ and define \tilde{r}_{v^*} in the analogous way. Observe that we have $\tilde{\ell}_{v^*} \leq \tilde{r}_{v^*}$ by the choice of p, q . Adding the resulting interval $[\tilde{\ell}_{v^*}, \tilde{r}_{v^*}]$ to J' yields an interval graph J on $V \cup V'$ that is consistent with both I and I' .

The second assertion of the lemma follows from the first. \square

The following observation is obvious:

Lemma 12. *If I is an interval graph on V as vertex set, then there is an interval graph J on V that is consistent with I and so that each of the intervals of J has its endpoints in $\{1, \dots, 2|V|\}$.*

We can now state the main result of this section.

Theorem 13. *There is an algorithm that, for any graph G with a given path decomposition of width w , determines in $2^{O(w^2 \log w)} \cdot |V(G)|$ -time a $d \in \mathbb{N}$ so that $d \leq \text{box}(G) \leq d + 1$.*




Figure 5: A box representation computed by the algorithm, where $d = 1$.

Proof. By omitting duplicated or empty bags, we may assume the given path decomposition $\mathcal{W} = \{W_1, \dots, W_t\}$ of G to have length $t \leq |V(G)|$. By Theorem 9, the boxicity of G is at most $w + 2$. Thus, it suffices to describe an algorithm that checks for some fixed $d \leq w + 2$, whether G has a box representation of dimension at most $d + 1$. This algorithm is then executed for every $d = 1, 2, \dots, w + 2$.

Our algorithm proceeds as follows.

1. Put $\mathcal{B}_0 = \{(\emptyset, \dots, \emptyset)\}$, which we consider as a tuple of d empty interval graphs.
2. For s from 1 to t do the following.
 - a. Compute all tuples of interval graphs (J_1, \dots, J_d) , where each J_i is an interval graph on the vertex set W_s so that all the interval endpoints are in $\{1, \dots, 2w + 2\}$, and where $G[W_s] = \bigcap_{i=1}^d J_i$.
 - b. For each of these tuples (J_1, \dots, J_d) , check whether there is a tuple $(J'_1, \dots, J'_d) \in \mathcal{B}_{s-1}$ that is consistent with (J_1, \dots, J_d) . If yes, add (J_1, \dots, J_d) to \mathcal{B}_s .
 - c. If $\mathcal{B}_s = \emptyset$ exit with $\text{box}(G) > d$.
3. Exit with $\text{box}(G) \leq d + 1$.

For the running time, note that the loop of line 2 is executed $t \leq |V(G)|$ times. Each execution of a–c requires $2^{O(w^2 \log w)}$ -time. Indeed, each interval graph J_i as in line 2a has vertex set W_s , and can thus be described by at most $2(w + 1)$ numbers, all of which are in $\{1, \dots, 2w + 2\}$. Thus, there are at most $(2w + 2)^{2w+2}$ possible such interval graphs J_i and therefore at most $(2w + 2)^{d(2w+2)}$ tuples considered in line 2a. Thus, each execution of lines 2a–c can be performed within the above claimed running time.

To verify that the algorithm is correct, we introduce a bit of notation and state two claims. Let us define an interval graph $I^* = ([\ell_v^*, r_v^*])_{v \in V(G)}$, where

$$\ell_v^* = \min_{v \in W_i} i \text{ and } r_v^* = \max_{v \in W_j} j.$$

We may perturb these points slightly, so that all endpoints become distinct. Note that I^* is the interval graph induced by the path decomposition of G ; that is $u, v \in V(G)$ are adjacent in I^* if and only if there is a bag W_s such that $u, v \in W_s$. Furthermore, we denote by G_s the induced subgraph $G[\bigcup_{s'=1}^s W_{s'}]$ of G on the first s bags of the path decomposition.

Inductively, we prove two claims:

for every $(J_1, \dots, J_d) \in \mathcal{B}_s$ there is a tuple (I_1, \dots, I_d) of interval graphs that is consistent with (J_1, \dots, J_d) and so that $G_s = \bigcap_{i=1}^d I_i \cap I^$.* (7)

and

if there are interval graphs I_1, \dots, I_d with $G_s = \bigcap_{i=1}^d I_i$ then there is a tuple $(J_1, \dots, J_d) \in \mathcal{B}_s$ that is consistent with (I_1, \dots, I_d) . (8)

Before proving the claims, we show how it follows from them that the algorithm is correct. Consider the case when the algorithm stops in line 3. Then $\mathcal{B}_t \neq \emptyset$, and with Claim (7) for $s = t$, we obtain $G = G_t$ as the intersection of $d + 1$ interval graphs, which proves that $\text{box}(G) \leq d + 1$. Now suppose the algorithm exits in line 2c, that is, that there is an s with $\mathcal{B}_s = \emptyset$. From Claim (8) we deduce that there is no tuple (I_1, \dots, I_d) of interval graphs with $G_s = \bigcap_{i=1}^d I_i$, which then precludes the existence of interval graphs I'_1, \dots, I'_d with $G = \bigcap_{i=1}^d I'_i$, as their restrictions $I'_i[\bigcup_{s'=1}^s W_{s'}]$ would intersect to G_s . Thus, $\text{box}(G) > d$, which finishes the proof of correctness.

We first prove Claim (7). For this, consider $(J_1, \dots, J_d) \in \mathcal{B}_s$. If $s = 1$ then $\bigcap_{i=1}^d J_i = G[W_1] = G_1$, and (7) is satisfied by setting $I_i = J_i$.

Hence let $s > 1$. By definition of \mathcal{B}_s there is a tuple $(J'_1, \dots, J'_d) \in \mathcal{B}_{s-1}$ that is consistent with (J_1, \dots, J_d) . Induction yields a tuple of interval graphs (I'_1, \dots, I'_d) that is consistent with (J'_1, \dots, J'_d) , and for which $G_{s-1} = \bigcap_{i=1}^d I'_i \cap I^*$. Applying Lemma 11 to each consistent pair I'_i and J_i yields an interval graph I_i consistent with both I'_i and J_i . Note that I_i is a supergraph of both I'_i and J_i . Since $G[W_s] = \bigcap_{i=1}^d J_i$ this means that G_s is a subgraph of $\bigcap_{i=1}^d I_i \cap I^*$.

Consider $i \in \{1, \dots, d\}$ and an edge e of I_i that is neither an edge of I'_i , nor of J_i . From Lemma 11 it follows that e has an endvertex u in $V(I'_i - J_i)$ and another endvertex v in $V(J_i - I'_i)$. Hence $u \in W_{s'} \setminus W_s$ for some $s' < s$ and $v \in W_s \setminus W_{s-1}$. In particular, the corresponding intervals, $[\ell_u^*, r_u^*]$ and $[\ell_v^*, r_v^*]$, of I^* do not intersect. Thus $I_i \cap I^* = (I'_i \cup J_i) \cap I^*$. Now, $G_{s-1} = \bigcap_{i=1}^d I'_i \cap I^*$ together with $G[W_s] = \bigcap_{i=1}^d J_i$ implies $G_s = \bigcap_{i=1}^d I_i \cap I^*$, as desired.

Finally, we show (8). If $s = 1$ then $\bigcap_{i=1}^d I_i = G[W_1]$, and thus, by Lemma 12, the algorithm shows in line 2a a tuple (J_1, \dots, J_d) of interval graphs that is consistent with (I_1, \dots, I_d) . Since every such tuple is consistent with $(\emptyset, \dots, \emptyset) \in \mathcal{B}_0$, it is added to \mathcal{B}_1 in line 2b.

Hence consider now $s > 1$. Letting I'_i be the restriction of I_i on $\bigcup_{s'=1}^{s-1} W_{s'}$, we see that induction yields a tuple $(J'_1, \dots, J'_d) \in \mathcal{B}_{s-1}$ that is consistent with (I'_1, \dots, I'_d) . Next, we apply Lemma 12 to $I_i[W_s]$ in order to obtain an interval graph J_i that is consistent with $I_i[W_s]$, and whose intervals have all their endpoints in $\{1, \dots, w + 2\}$. Consequently, the tuple (J_1, \dots, J_d) is among the

tuples computed in step 2a of the algorithm. Moreover, (J_1, \dots, J_d) is consistent with (I'_1, \dots, I'_d) , and thus, also consistent with $(J'_1, \dots, J'_d) \in \mathcal{B}_{s-1}$. Therefore, (J_1, \dots, J_d) is added to \mathcal{B}_s in step 2b of the algorithm. \square

We mention that, while the algorithm as given only computes the number d , we can also recover a concrete box representation of dimension $d + 1$. For this, it suffices to store for each tuple in \mathcal{B}_s to which tuple in \mathcal{B}_{s-1} it is consistent (if there are more, we simply choose one).

Together with the algorithm of Bodlaender [6] that computes a path-decomposition of a graph G of width $\text{pw}(G)$ in $f(\text{pw}(G)) \cdot |V(G)|$ time, we obtain Theorem 2. We note that the running time could conceivably be improved by using a faster approximation algorithm with, say, a constant approximation factor.

5 Proof of Theorem 3

It is an open problem whether boxicity is polynomial-time solvable on graphs of bounded treewidth. While we cannot solve the problem, we can offer an indication why we suspect boxicity to be hard.

The first approach to prove tractability is usually dynamic programming. Evidently, this is because Courcelle [8] proved that a vast number of problems, namely those expressible in monadic second order logic, can be solved in polynomial time by a generic dynamic programming algorithm, if the treewidth is bounded. However, nobody appears to know how to formulate “ $\text{box}(G) \leq d$?” in monadic second order logic, and it is doubtful that this is possible at all. More generally, dynamic programming seems to fail. Why is that so? We think this is because the tree-like structure of the input graph does not translate to a tree-like structure in the interval representation: given an input graph G of bounded treewidth, it may very well be the case that at least one interval graph in any optimal interval representation of G has unbounded treewidth.

To illustrate this, consider a $K_{2,n}$, where the smaller bipartition class is comprised of two vertices x and y , and the larger consists of v_1, \dots, v_n . Clearly, $K_{2,n}$ has pathwidth 2 and boxicity 2 as well: in fact, $K_{2,n} + xy$ and $K_{2,n} + \{v_i v_j : i, j\}$ are two interval graphs whose intersection is $K_{2,n}$. Now, let I_1, I_2 be any two interval graphs with $K_{2,n} = I_1 \cap I_2$. The vertices x and y are not adjacent in at least one of I_1 and I_2 , say in I_1 . Suppose that I_1 contains a pair of non-adjacent v_i, v_j : then $xv_i y v_j x$ is an induced 4-cycle, which is impossible in an interval graph. Thus, $\{v_i\}_{i=1}^n$ form a clique of size n in I_1 , and I_1 has therefore pathwidth at least $n - 1$.

What about stronger width-parameters? We have found a similar, albeit more complicated, example for bounded bandwidth, a parameter even more restrictive than pathwidth. Theorem 3 is a direct consequence of the following lemma.

Lemma 14. *For every n there is a graph G^n of bandwidth at most 16 and boxicity 2, so that in any interval representation $G = I_1 \cap I_2$ one of I_1 and I_2 has treewidth $\geq |V(G^n)|/32$.*

In light of the lemma, we would like to strengthen the conjecture of Adiga et al. [3]: We believe that BOXICITY remains NP-complete even for graphs of

bounded bandwidth.

Proof of Lemma 14. As a basic building block for G^n we use copies of the graph B , which consists of a path $w_1 \dots w_6$ and two vertices u, v adjacent to each of w_1, \dots, w_6 but not to each other; see Figure 6. Clearly, B has boxicity 2, and moreover, if B is represented as the intersection of two interval graphs I_1, I_2 then

$$\text{for some } k \in \{1, 2\}, uv \in E(I_k) \text{ and } w_1, \dots, w_6 \text{ is a clique in } I_{3-k}. \quad (9)$$

This follows directly from the fact that $uw_i v w_j u$ is an induced 4-cycle in B for each $1 \leq i < j - 1 < 6$.

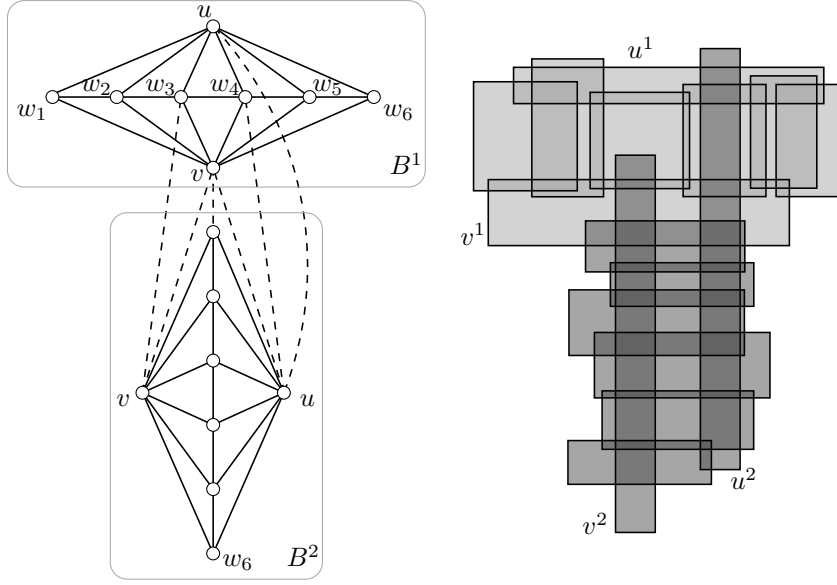


Figure 6: Glueing the gadgets (left) and a geometric realisation (right)

We define graphs G^1, G^2, \dots, G^n iteratively by taking as G^1 a copy B^1 of B with vertex set $u^1, v^1, w_1^1, \dots, w_6^1$. Then, given G^i we obtain G^{i+1} by adding another copy B^{i+1} of B on vertex set $u^{i+1}, v^{i+1}, w_1^{i+1}, \dots, w_6^{i+1}$, where we make

- u^{i+1} adjacent to u^i, v^i and w_4^i ;
- v^{i+1} adjacent to v^i and w_3^i ; and
- w_1^{i+1} adjacent to v^i .

An indication that G^n has indeed boxicity 2 as claimed is given in Figure 6. That G^n has bandwidth ≤ 16 can also easily be checked. We fix two interval graphs I_1, I_2 , so that $G^n = I_1 \cap I_2$.

First, we prove that the edge uv flips between consecutive copies of B , that is

$$\text{for some } k \in \{1, 2\}, u^i v^i \in E(I_k) \text{ implies } u^{i+1} v^{i+1} \in E(I_{3-k}). \quad (10)$$

To keep notation simple, we consider the case when $i = 1$, and we assume that $u^1 v^1 \in E(I_1)$, so that our task is to show that $u^2 v^2 \in E(I_2)$.

Observe that for each $j = 1, 2, 3, 5, 6$ the vertices $u^1 w_j^1 v^1 u^2 u^1$ form a 4-cycle in G^n and thus in I_2 . Since I_2 is chordal but u^1 and v^1 are not adjacent in I_2 , it follows that $u^2 w_j^1 \in E(I_2)$. As this edge is not present in G^n , we have consequently that

$$w_4^1 \text{ is the only neighbour in } I_1 \text{ of } u^2 \text{ among } w_1^1, \dots, w_6^1. \quad (11)$$

The graph (a) in Fig. 2 we call a *Y*, the graph (b) we call an *umbrella*. As both I_1 and I_2 are interval graphs, neither of them contains a *Y* or an *umbrella* as an induced subgraph.

We next show that

$$v^1 w_2^2 \in E(I_1). \quad (12)$$

From (11), we deduce that in I_1 the induced path $w_2^1 \dots w_6^1$ together with u^2 and w_2^2 forms an induced *Y*, unless w_2^2 is adjacent to at least one of the vertices w_2^1, \dots, w_6^1 . Thus, we can choose j be such w_j^1 is a neighbour of w_2^2 in I_1 and closest to w_4^1 . If $j \neq 4$ then $w_j^1 \dots w_4^1 u^2 w_2^2$ forms an induced cycle of length ≥ 4 , which contradicts the chordality of I_1 . Thus, $w_2^2 w_4^1$ is an edge of I_1 .

Now $w_2^1 \dots w_6^1$ together with v^1 and w_2^2 form an induced umbrella in I_1 , unless w_2^2 has further neighbours among w_2^1, \dots, w_6^1 , or $w_2^2 v^1 \in E(I_1)$. In the latter case, we have proved (12), so assume that in I_1 , w_2^2 has a second neighbour in $\{w_2^1, \dots, w_6^1\}$. As I_1 is chordal one of w_3^1 and w_5^1 has to be adjacent to w_2^2 . By symmetry, let us say that $w_3^1 w_2^2 \in E(I_1)$. Then $w_3^1 w_2^2 u^2 v^1 w_3^1$ is a 4-cycle that is induced unless w_2^2 and v^1 are adjacent in I_1 (here we use (11) again). This proves (12).

Turning back to G^n , we observe that $v^1 u^2 w_2^2 v^2 v^1$ form an induced 4-cycle. Thus, each of I_1 and I_2 must contain exactly one of the diagonals $v^1 w_2^2$ and $u^2 v^2$. Since we have already proved that $v^1 w_2^2 \in E(I_1)$ it follows that $u^2 v^2 \in E(I_2)$. This finishes the proof of (10).

Next, we see that consecutive copies of B rotate by 90° . More formally, we associate with a gadget B^i a matrix $D^i \in \mathbb{R}^{2 \times 2}$, where the first row encodes the orientation of the induced path $w_1^i \dots w_6^i$, while the second row corresponds to u, v . If $u^i v^i \in E(I_k)$ for $k = 1$ or $k = 2$, we set

$$D_{1k}^i = \begin{cases} 1 & \text{if } \ell_k(w_1^i) < \ell_k(w_6^i) \\ -1 & \text{otherwise} \end{cases} \quad \text{and } D_{1,3-k}^i = 0$$

and

$$D_{2k}^i = 0 \text{ and } D_{2,3-k}^i = \begin{cases} 1 & \text{if } r_{3-k}(v^i) < \ell_{3-k}(u^i) \\ -1 & \text{otherwise} \end{cases}$$

By symmetry, we may assume B^1 to be embedded in such a way that $u^1 v^1 \in E(I_1)$, $r_2(v^1) < \ell_2(u^1)$ and $\ell_1(w_1^1) < \ell_1(w_6^1)$. Then D^1 is simply the identity matrix. Below we see that D^i is the $(i-1)$ th power of the rotation matrix $R := \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$.

We only prove $D^2 = R$. For larger i this follows with analogous arguments. As we assumed $u^1 v^1 \in E(I_1)$, it follows from by (10) that $u^2 v^2 \in E(I_2)$, which implies $D_{11}^2 = 0$ and $D_{22}^2 = 0$. Let us next show that $D_{21}^2 = 1$, that is

$$r_1(v^2) < \ell_1(u^2). \quad (13)$$

Indeed, as $u^2v^2 \in E(I_2)$, we either have $r_1(v^2) < \ell_1(u^2)$ or $r_1(u^2) < \ell_1(v^2)$. By (9), $w_1^1 \dots w_6^1$ is an induced path in I_1 , and from the assumption that $\ell_1(w_1^1) < \ell_1(w_6^1)$ we deduce $\ell_1(w_3^1) < \ell_1(w_4^1)$. From (11) it follows that $\ell_1(u^2) \in (r_1(w_3^1), r_1(w_4^1)]$. Since v^2 is adjacent to w_3^1 we now see that $r_1(u^2) < \ell_1(v^2)$ is impossible as this would imply $r_1(w_3^1) < r_1(u^2) < \ell_1(v^2)$. This shows (13).

To show the rotation, it remains to prove that

$$\ell_2(w_1^2) > \ell_2(w_6^2), \quad (14)$$

which is to say that $D_{12}^2 = -1$.

Suppose that $r_2(v^1) \leq r_2(w_1^2)$. Then since $r_2(v^1) < \ell_2(u^1)$ it follows that each of w_1^1, \dots, w_6^1 is adjacent to w_1^2 in I_2 , from which we deduce that w_1^2 is not a neighbour of any of w_1^1, \dots, w_6^1 in I_1 . However, by (11), then $w_2^1 \dots w_6^1$ induces with u^2 and w_1^2 a Y -subgraph in I_1 , which is impossible. Thus $r_2(w_1^2) < r_2(v^1)$.

Now, if $\ell_2(w_1^2) < \ell_2(w_6^2)$ then one of w_2^2, \dots, w_6^2 must be adjacent to v^1 in I_2 . Since I_2 is a chordal graph and $w_1^2 \dots w_6^2$ an induced path in I_2 , this is only possible if $v^1w_2^2$ is an edge of I_2 . But, by (12), we also have $v^1w_2^2 \in E(I_1)$, which contradicts $v^1w_2^2 \notin E(G^n)$. Therefore, we have proved (14).




Figure 7: Spiral of v^i -boxes, and vertical line intersecting $\geq n/4$ boxes

Summing up, the boxes corresponding to the induced path $v^1v^2 \dots v^n$ are arranged in a spiral pattern. This spiral has two options: either it may become ever larger or, after a number of steps, it may become smaller and smaller. In both cases, we find a vertical line that meets a quarter of the v^i -boxes, which translates to a clique of size $n/4$ in I_1 . In particular, $\text{tw}(I_1) \geq n/4$. As $|V(G)| = 8n$, the proof is complete. \square

6 Discussion

In this paper, we treated BOXICITY from the perspective of parameterized complexity. We presented

- a parameterized algorithm for BOXICITY with respect to the parameter cluster vertex deletion;
- an additive 1-approximation algorithm for BOXICITY that runs in $2^{O(w^2 \log w)} \cdot n$ time where w is the width of a given path decomposition of the input graph;

- and a family of graphs of bounded bandwidth that need, in any optimal box representation, dimensions of unbounded treewidth.

In some respect, the method of our first algorithm is a generalization of the true twin reduction. The key insight is that if there are many vertex sets (the clusters) that are identical in the graph then many of these sets will have essentially the same geometric realization. Deleting one of these many “geometric twins” is unlikely to change boxicity.

We believe this approach can be exploited further. Indeed, we are convinced that with similar methods as developed in this article, we can also formulate a parameterized algorithm for BOXICITY when the parameter is *distance to stars* – the smallest number of vertices whose removal results in a disjoint union of stars. Like cluster vertex deletion, distance to stars provides a non-trivial parameterization for BOXICITY between vertex cover (solved) and feedback vertex set (open). Moreover, given a graph G , computing a minimum set $X \subseteq V(G)$ such that $G[V - X]$ is a disjoint union of stars can be done in $f(|X|) \cdot |V(G)|^{O(1)}$ time [17].

Our second algorithm yields an additive 1-approximation for BOXICITY on graphs of bounded pathwidth. Two questions that immediately arise are: can we get rid of the additive 1, such that the algorithm computes $\text{box}(G)$ exactly? Can the algorithm be lifted to run on graphs of bounded treewidth?

We stated earlier our impression of the first question – we conjecture BOXICITY to be NP-complete on graphs of bounded bandwidth, thus including graphs of bounded pathwidth. Our reasoning is that, in any optimal representation as the intersection of interval graphs, some of the interval graphs may have unbounded treewidth, even if the input graph has bounded bandwidth. This seems to annul the main advantage of bounded pathwidth, that the number of box representation of the union of previous bags can be compressed to a size bounded by a function of the pathwidth.

We turn to the second question: why is it difficult to extend the algorithm to graphs of bounded treewidth? We rely heavily on the fact that the one extra dimension is sufficient to reflect the path decomposition of the whole graph. If we mimick this approach for bounded treewidth we have to describe the tree decomposition of the graph with as few extra dimensions as possible. How many extra dimensions would we need? As many as the boxicity of the chordal supergraph obtained by turning each bag of the decomposition into a clique. If we started with a path decomposition, the boxicity will be one. For a general tree decomposition, however, it could well be that the boxicity of this chordal graph is about the treewidth of the input graph [7]. This suggests that there might be input graphs G for which $\text{box}(G)$ is much lower than the number of dimensions required to describe their tree decomposition, which makes it impossible to approximate using only the techniques of Section 4.

References

- [1] A. Adiga, J. Babu, and L. S. Chandran. Polynomial time and parameterized approximation algorithms for boxicity. In *Proceedings of the 7th International Symposium on Algorithms and Computation (IPEC 2012)*, LNCS 7535, pages 135–146, 2012.

- [2] A. Adiga, D. Bhowmick, and L. S. Chandran. The hardness of approximating the boxicity, cubicity and threshold dimension of a graph. *Discrete Applied Mathematics*, 158(16):1719–1726, 2010.
- [3] A. Adiga, R. Chitnis, and S. Saurabh. Parameterized algorithms for boxicity. In *Proceedings of the 21st International Symposium on Algorithms and Computation (ISAAC 2010)*, LNCS 6506, pages 366–377, 2010.
- [4] E. Asplund and B. Grünbaum. On a coloring problem. *Math. Scand.*, 8:181–188, 1960.
- [5] A. Bielecki. Problem 56. *Colloq. Math.*, 1:333, 1948.
- [6] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305–1317, 1996.
- [7] L. S. Chandran and N. Sivadasan. Boxicity and treewidth. *Journal of Combinatorial Theory, Series B*, 97(5):733–744, 2007.
- [8] B. Courcelle. The monadic second-order logic of graphs I. Recognizable sets of Finite Graphs. *Information and Computation*, pages 12–75, 1990.
- [9] M. Cozzens. Higher and multi-dimensional analogues of interval graphs. Ph.d. thesis, Department of Mathematics, Rutgers University, New Brunswick, NJ, 1981.
- [10] R. Diestel. *Graph Theory* (4th edition). Springer-Verlag, 2010.
- [11] M. Doucha and J. Kratochvíl. Cluster vertex deletion: A parameterization between vertex cover and clique-width. In *Proceedings of the 37th International Symposium on Mathematical Foundations of Computer Science (MFCS 2012)*, LNCS 7464, pages 348–359, 2012.
- [12] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Springer-Verlag, 2013.
- [13] L. Esperet. Boxicity of graphs with bounded degree. *European Journal of Combinatorics*, 30(5):1277–1280, 2009.
- [14] R. Ganian. Twin-cover: Beyond vertex cover in parameterized algorithms. In *Proceedings of the 6th International Symposium on Algorithms and Computation (IPEC 2011)*, LNCS 7112, pages 259–271, 2011.
- [15] A. Kostochka. Coloring intersection graphs of geometric figures with a given clique number. In J. Pach, editor, *Towards a theory of geometric graphs*, volume 342 of *Contemp. Math.*, pages 127–138. Amer. Math. Soc., 2004.
- [16] J. Kratochvíl. A special planar satisfiability problem and a consequence of its NP-completeness. *Discrete Applied Mathematics*, 52(3):233–252, 1994.
- [17] N. Nishimura, P. Ragde, and D. M. Thilikos. Fast fixed-parameter tractable algorithms for nontrivial generalizations of vertex cover. *Discrete Applied Mathematics*, 152:229–245, 2005.

- [18] F. S. Roberts. On the boxicity and cubicity of a graph. *Recent Progress in Combinatorics*, pages 301–310, 1969.
- [19] E. Scheinerman. Intersection classes and multiple intersection parameters. Ph.d. thesis, Princeton University, 1984.
- [20] J. Spinrad. *Efficient Graph Representations*. Fields Institute monographs. American Mathematical Society, 2003.
- [21] C. Thomassen. Interval representations of planar graphs. *Journal of Combinatorial Theory, Series B*, 40(1):9–20, 1986.
- [22] M. Yannakakis. The Complexity of the Partial Order Dimension Problem. *SIAM Journal on Algebraic and Discrete Methods*, 3(3):351–358, 1982.

Version February 21, 2014

Henning Bruhn <henning.bruhn@uni-ulm.de>
 Morgan Chopin <morgan.chopin@uni-ulm.de>
 Felix Joos <felix.joos@uni-ulm.de>
 Institut für Optimierung und Operations Research
 Universität Ulm, Germany
 Oliver Schaudt <schaudt@uni-koeln.de>
 Institut für Informatik
 Universität zu Köln, Germany